TRec: Relay-Resilient Transparent Authentication using Trajectory Recognition

Mika Juuti*, Christian Vaas†, Ivo Sluganovic†, Hans Liljestrand‡, N. Asokan‡ and Ivan Martinovic†
*Aalto University, mika.juuti@aalto.fi
†University of Oxford, christian.vaas@cs.ox.ac.uk, ivo.sluganovic@cs.ox.ac.uk, ivan.martinovic@cs.ox.ac.uk
‡Aalto University and University of Helsinki, hans.liljestrand@aalto.fi, asokan@acm.org

arXiv:1610.02801v1 [cs.CR] 10 Oct 2016

***Abstract*—Transparent authentication (TA) schemes are those in which a user is authenticated by a verifier without requiring explicit user interaction. Those schemes promise high usability and security simultaneously. Many TA schemes rely on the received signal strength as an indicator for the proximity of a user device (prover). However, such implicit proximity verification is not secure against an adversary who can relay messages.**

In this paper, we propose a novel approach for thwarting relay attacks in TA schemes: the prover permits access to authentication credentials only if it can confirm that it is near the verifier. We present TRec, a system for relay-resilient transparent authentication in which the prover does proximity verification by identifying its approach trajectory to the intended verifier and comparing it with known authorized reference trajectories. Trajectories are measured using low-cost sensors commonly available on personal devices. We demonstrate the security of TRec against a class of adversaries and its ease-of-use by analyzing empirical data, collected using a TRec prototype. TRec is efficient and can be easily integrated into existing TA schemes.

I. Introduction

Authentication of users is necessary for access control to data or physical objects. The predominant approach to user authentication is to use passwords. This is widely acknowledged to fall short in terms of both usability and security [3]. Effective alternatives to password-based authentication are yet to emerge [2]. This has sparked the quest for *transparent authentication*. Its aim is to improve security through user acceptance by not requiring explicit user interaction solely for authentication. Instead, transparent authentication schemes exploit cues like behavior [32] biometrics [26] or environmental context [29]. One class of transparent authentication schemes, known as "zero-interaction authentication" (ZIA) [5], rely on *verifier* $\mathcal{V}$ authenticating a user when a *prover* device $\mathcal{P}$ carried by the user is nearby. In ZIA schemes $\mathcal{V}$ verifies proximity of $\mathcal{P}$ by measuring the strength of radio signals emitted from $\mathcal{P}$ over some short range wireless channel. Such authentication schemes are widely deployed, for instance in "keyless entry and start" systems in cars or in open source projects like BlueProximity [25]. But they are vulnerable to *relay attacks* [9], [10], where the attacker relays messages between $\mathcal{P}$ and $\mathcal{V}$ when they are not co-located, leading $\mathcal{V}$ to erroneously conclude that $\mathcal{P}$ is nearby.

Known defenses to these attacks range from distance bounding protocols [4] to comparing the ambient context between $\mathcal{P}$ and $\mathcal{V}$ [11], [29]. However, they face deployment challenges. For example, context-based relay attack resistance has unclear security guarantees [24]. Distance bounding, on the other hand, requires precise timing: 1 ns measurement error affects the estimated distance already by 15 cm. This implies requiring additional hardware and low-level software changes.

In this paper we propose a novel approach to thwart relay attacks against proximity-based transparent authentication systems by having $\mathcal{P}$ verify proximity to an intended $\mathcal{V}$ before *allowing access* to the credentials used in the authentication protocol. We present TRec, a system that enforces such proximity verification by $\mathcal{P}$ using on-board microelectromechanical system (MEMS) sensors to measure its *approach trajectory* towards the intended $\mathcal{V}$ and comparing it with known authorized reference trajectories. A central design principle in TRec is to rely only on sensors that monitor $\mathcal{P}$'s own movement (e.g. accelerometer and gyroscope) rather than on sensors that measure environmental factors (e.g. location, radio signal emission, or ambient properties) that can be manipulated by an adversary from the outside. We built TRec as an Android application and used it to gather trajectory data of 20 different routes in two cities (totaling 123 km). Using this dataset, we show that TRec has acceptable false reject (FRR) and false accept (FAR) rates.

Commodity devices use low-cost MEMS sensors that are noisy, include bias terms and miss data. Designing TRec to work on commodity devices raised several technical challenges, leading to questions like "how to effectively represent a trajectory using accelerometer/gyroscope measurements only?" and "what is the best way to compare two trajectories?", while considering energy budgets of portable devices. We identify and address these challenges and systematically evaluate the resulting system. Briefly, our contributions are the following:

- We propose using **prover-side proximity verification to resist relay attacks** against proximity-based transparent authentication systems (Sections II and III).
- We design and implement a **concrete system, TRec**, incorporating this idea by addressing several challenges in measuring prover's approach trajectory and using it to determine proximity to the verifier (Section III).
- By systematically analyzing trajectory data in two cities, we **demonstrate the security and usability of TRec** (Section IV). We also show that TRec's average **energy consumption is low**: we estimate that under typical usage conditions, the battery drain due to TRec over the course of a work day is in the range of 4%-7% of battery capacity. (Section IV-F).

II. Concept and Assumptions

A. System model

Figure 1 shows a model for proximity-based authentication. Its goal is to let a verifier $\mathcal{V}$ convince itself that a user $\mathcal{U}$ is nearby. $\mathcal{U}$ carries a personal device $\mathcal{P}$. Authentication itself is based on a challenge-response protocol using a previously established security association, e.g. a shared symmetric key between $\mathcal{P}$ and $\mathcal{V}$. In addition to verifying authenticity of $\mathcal{P}$, $\mathcal{V}$ also verifies its proximity to $\mathcal{P}$.

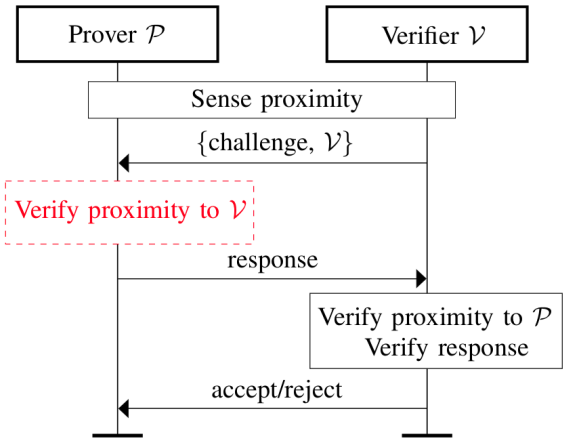

Fig. 1. Transparent authentication: challenge-response protocol triggered upon sensing proximity of $\mathcal{P}$ to $\mathcal{V}$. The proximity verification component introduced in our approach is shown as a red/dashed box.

To protect against relay attacks, we augment this model by having $\mathcal{P}$ regulate access to the authentication credentials by first verifying its proximity to $\mathcal{V}$. In particular, we propose that $\mathcal{P}$ does this proximity verification by measuring its approach trajectory towards $\mathcal{V}$ (Section III-C). If proximity verification fails, $\mathcal{U}$ is asked to explicitly confirm whether he is near $\mathcal{V}$.

B. Adversary Model

We assume an adversary $\mathcal{A}$ who has deployed a wireless relay providing it with Dolev-Yao [6] capabilities. Although $\mathcal{A}$ can control the message flow between $\mathcal{P}$ and $\mathcal{V}$, it cannot break the cryptographic protection of a secured channel. Nevertheless, by relaying the challenge and response $\mathcal{A}$ can successfully bypass the authentication mechanism as well as the proximity verification by $\mathcal{V}$ where $\mathcal{V}$ will be measuring $\mathcal{A}$'s signal strength rather than $\mathcal{P}$'s. We also assume that $\mathcal{A}$ can manipulate ambient properties. For instance, it can relay audio so that $\mathcal{P}$ and $\mathcal{V}$ perceive the same ambient audio [24]. We target an adversary who follows the legitimate user. We do not address an adversary who can steal or otherwise obtain physical access to the prover device $\mathcal{P}$. Local user authentication techniques, such as biometric continuous authentication [21] can resist such adversaries.

C. Design Goals and Challenges

Goals. We set the following goals for our relay-resilient proximity verification system:

R1. **Usability**: Transparent authentication must minimize explicit user action. If trajectory comparison fails when $\mathcal{U}$ is in fact near $\mathcal{V}$, $\mathcal{U}$ will be required to fall back to explicit proximity confirmation. Our system should therefore minimize the false reject rate.

R2. **Security**: The system should not incorrectly conclude that $\mathcal{U}$ is near $\mathcal{V}$ even in the presence of a relay. Therefore it should minimize the false accept rate.

R3. **Efficiency**: The computational and energy costs of proximity verification should be insignificant.

R4. **No external signals**: Since $\mathcal{A}$ can control ambient properties, proximity verification should not depend on any external signals.

R5. **Local decision-making**: Proximity verification must be carried out entirely within $\mathcal{P}$.

There are two rationales for R5. One is **privacy**: data collected for proximity verification should not be exposed to any third party. The other is **deployability**: a local solution can be seamlessly integrated into any proximity-based transparent authentication scheme by only modifying $\mathcal{P}$ without having to change the actual protocol.

Challenges. Designing a system to meet the above goals raises several challenges. First, we need to understand what kind of trajectory characteristics can be measured repeatedly reliably using commodity MEMS sensors (Section III-A). Second, using noisy sensors to measure trajectory necessarily requires a balance between tolerating innocuous errors to minimize false rejects while avoiding false accepts. Innocuous measurement errors can be reduced by filtering and smoothing (Section III-B). Third, multiple instances of a trajectory can improve trajectory comparison accuracy. However, traversing trajectories is time consuming; users should ideally be able to use the system with little or no training delay (Section III-D). Fourth, MEMS sensor recordings may include gaps: mobile OSs save energy by suspending the CPU when there is no foreground activity, which can result in sensor hardware buffers overflowing, leading to data loss. Long sensory gaps can be devastating for performance (Section III-E).

III. Architecture and Implementation

We now describe `TRec`, our system that uses prover-side proximity verification to prevent relay attacks.

A. Trajectory Representation

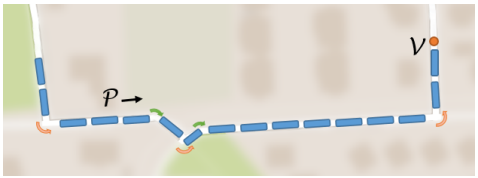

Fig. 2. $\mathcal{P}$'s approach trajectory towards $\mathcal{V}$ is described by a set of primitive steps needed to reach %!TEX root = ../main.texver. In this example, the steps are *move 2, left 90°, move 3, right 30°, move 1, left 90°, move 1, right 30°, move 7, left 90°, move 2.*

To satisfy requirement R4, we avoid external, insecure, data sources like GPS [27] or ambient sensor modalities and rely only on *gyroscope* and *accelerometer* to capture user's movement. We represent a trajectory as a *temporally ordered*

sequence of discrete *primitives* consisting of segments of *movement* interleaved with turns (*left* or *right*) derived from angular information. (See Figure 2 for an example.)

An intuitive way to represent a trajectory is as a sequence of coordinates, like in dead reckoning [16]. However, a one-dimensional sequence of primitives is more robust to sensor noise than a two-(or even three-) dimensional coordinate: the impact of a missed turn on the resulting sequence is less than it is on the outcome of dead reckoning.

Using sensor data, we recognize two streams of primitives: $(M/S, t_i)$ symbols (for "movement" or "stationary" at time t_i) are generated at a fixed rate and $(L/R, t_i)$ symbols (turn "left" or "right" at time t_i) are generated opportunistically whenever a turn is detected. The two streams are then combined into one sequence, with turn events taking precedence.

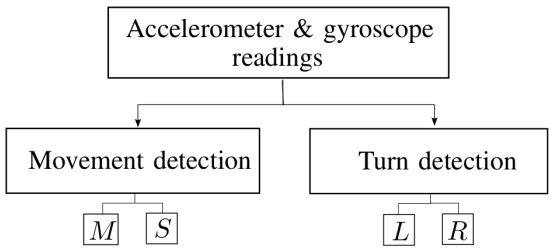

Fig. 3. Primitive generation: movement M, stationary S, turns L, R.

The overall system has three essential parts: primitive generation, trajectory comparison and updates.

B. Primitive Generation

Turn primitives. By experimenting with different sampling rates and sensors, we determined that a 20 Hz sampling rate is sufficient to detect turns with a granularity of 15°. We project the gyroscope to ground direction, obtain the heading angle by integrating the angular speed and then detect turns when the 2s sliding window standard deviation of the heading angle is above a threshold ($\sigma_1 = 3°$). However, sudden gravity shifts[1] cause errors in turn estimation: we disregard gyroscope data as unreliable in such situations. To remove drift in MEMS gyroscopes, we use a high-pass filter, where gyroscope measurements smaller than 8.6°/s (0.15 rad/s) are exponentially weighted down. Fine-grained beginnings and ends of turns are found where the sliding window standard deviation is above a smaller threshold ($\sigma_2 = 1°$). The turn detection system assumes that $\mathcal{P}$ has reliable gravity estimates: $\mathcal{P}$ could for instance be integrated into a vehicle or firmly attached to the body to avoid collusion of gravity direction. In this paper we collect data by integrating $\mathcal{P}$ with a bicycle.

Movement primitives. To identify movement, we use a logistic regression (LR) algorithm [19] that continuously predicts movement mode at one second intervals. The prediction is done on-the-spot, and does not take previous prediction results into account. In reality however, two successive events are dependent. We additionally use a Hidden Markov Model (HMM) [7] to capture this dependency.

In HMMs, probabilities to move between hidden states are modeled with a first-order Markov chain (*transition probabilities*). Each hidden state gives a cue about itself by emitting a primitive at any given time. In our case, we are trying to determine movement primitives (M or S) by observing the output of LR. We use HMM Viterbi algorithm [7] to smoothen the observations into the most likely sequence of primitives. Finally, we determine the representative primitive for each five measurements as the most frequent among the five. This is done in order to reduce computational complexity during comparison (Section III-C). This scheme gives us a continuous stream of one M or S primitive every five seconds.

C. Trajectory Comparison

We differentiate between two types of trajectories. *Reference paths* are previously authorized trajectories of a $\mathcal{P}$ towards some verifier $\mathcal{V}$. *Candidate paths* are trajectories towards $\mathcal{V}$ perceived by $\mathcal{P}$ before a proximity authentication session. These can contain errors introduced through noisy measurements. To verify proximity to $\mathcal{V}$, $\mathcal{P}$ compares the candidate path to the reference path. Figure 4 shows our proposed proximity verification scheme. If the trajectory comparison succeeds, $\mathcal{P}$ continues with the authentication protocol by computing the response to the challenge and sending it to $\mathcal{V}$. If trajectory comparison fails, the system falls back to explicit trajectory confirmation by $\mathcal{U}$. Successful explicit confirmation implies that $\mathcal{P}$'s candidate path can be added to the trajectory repository as an authorized trajectory: a reference path.

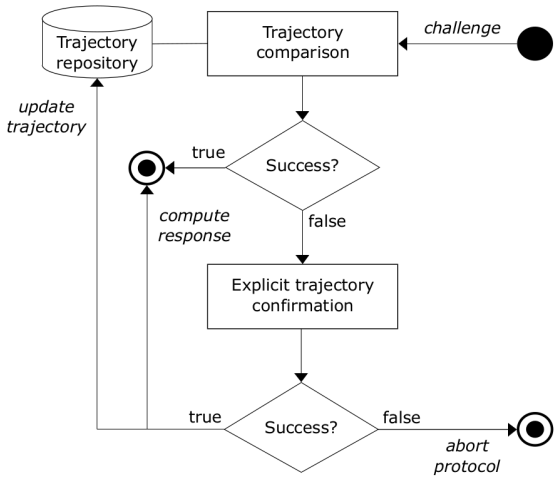

Fig. 4. Prover-side proximity verification. $\mathcal{P}$ compares the current candidate path against reference paths in the repository. $\mathcal{P}$ computes the response to $\mathcal{V}$'s challenge if and only if candidate path matches a reference path.

Candidate and reference paths are represented as sequences of characters as discussed in Section III-A. Trajectory comparison is therefore a similarity comparison between strings. We evaluated several string matching metrics and chose to the use Needleman-Wunsch (NW) similarity[2] [7], which can be viewed as a combination of the longest common subsequence and edit distance algorithms. In NW terminology, both insertions and deletions are called gaps. We chose the

[1]Gravity in Android is a software sensor (a low-pass filter on raw accelerometer data) which takes a few seconds after an orientation change to stabilize.

[2]Needleman-Wunsch had the best FAR/FRR trade-offs among the tested algorithms on our dataset in Section IV.

simple parametrization: match *+1*, mismatch *-2* and gap *-1*. Matches increment the score slightly, while mismatches are always worth one insert and delete.Matches can be seen as evidence and mismatches counter-evidence that two sequences are related. Before comparison, S symbols are removed and timestamps (t_i) are used to trim strings to the same temporal length.

Instances of the same reference path will differ due to noise from various sources. Therefore, we need to establish a *decision threshold* that determines how much noise is acceptable. If the similarity score is higher than the threshold, the candidate path is accepted (that is, *recognized* to be the same as the reference path). Otherwise it is rejected. This introduces a trade-off between usability (FRR) and security (FAR). An *initial threshold* that has a good FAR/FRR trade-off is determined prior to deployment (See Section IV-C).

D. Updating Reference Paths

Once the system is deployed, we use feedback from failed and successful trajectory recognition attempts to adjust the decision threshold. The initial threshold might under- or overestimate the variation in future instances of a given reference path r. The decision threshold should be adjusted to achieve a better FAR/FRR trade-off, either by decreasing (better usability) or increasing (better security) it.

We call such a path-specific decision threshold a *local threshold*. To compute local thresholds for a given r we need instances i_r of r and instances $i_{\bar{r}}$ of reference paths towards other verifiers $\bar{\mathcal{V}}$. These are used to calculate *within-* and *between-class similarities*. The better the separability of these two types of similarities, the better FAR/FRR trade-offs we can achieve.When a user $\mathcal{U}$ starts using `TRec` we do not have enough instances i_r, $i_{\bar{r}}$ to compute within- and between-class similarities. Instances i_r will be gradually collected as the user repeatedly traverses r. We propose two ways to acquire instances $i_{\bar{r}}$ of paths towards another $\bar{\mathcal{V}}$:

- use trajectories generated from a map for the current geographic region, or
- collect all trajectories of a given user and create a generative probabilistic model (Markov chain) to simulate new reference paths instances.

We choose the latter option in `TRec`. In the beginning, `TRec` will only observe a few instances of a reference path and any new local decision threshold we obtain would be severely over-learned. Naively trusting the seen instances risks a small sample size fallacy. Therefore, we need a way to model the trustworthiness of the estimated local threshold. We model the confidence as a mixture model using a convex combination [19] of the thresholds d_i (initial) and d_l (local) with a confidence parameter $\lambda \in [0, 1]$:

$$d = \lambda d_l + (1 - \lambda) d_i. \tag{1}$$

We call the resulting threshold d a *mixed threshold*. A common way to model the confidence in small sample sizes is to use add-one smoothing [19]. When n is the number of seen instances of a reference path, we can model λ as:

$$\lambda(n) = \frac{n - 1}{n}. \tag{2}$$

This fulfills our boundary conditions for λ: $\lambda(1) = 0$, implies no confidence when we have only seen one instance of a reference path and $\lambda(\infty) = 1$, signifying full confidence with infinitely many instances. Figure 5 shows how the confidence factor increases w.r.t. the number of reference path instances seen so far. We use equations 1 and 2 to determine mixed thresholds throughout this paper. The mixed, local and initial thresholds are evaluated in Section IV-E.

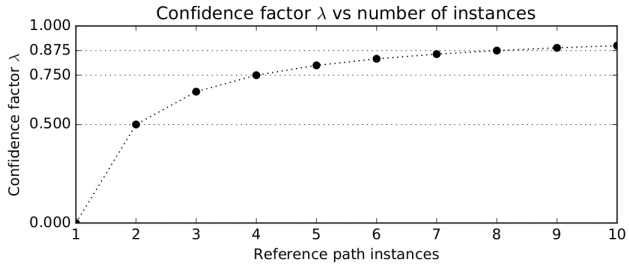

Fig. 5. The confidence factor λ increases with the number of seen reference path instances. The distance to one halves at every multiple of two.

E. Android Prototype for `TRec`

Figure 6 shows the UI of the `TRec` Android app. To reduce battery drain, the device enters sleep after experiencing 5 minutes of non-movement. `TRec` acquires a wakelock [1] when significant motion is detected to ensure that no relevant sensor data is lost due to sleep.

Fig. 6. `TRec` classifies movement at one second intervals. *Left:* Users can quickly lock $\mathcal{V}$ in case of false accepts. *Right:* Interface for saving unrecognized paths. The length of the reference path can be adjusted.

Resource requirements. `TRec` expects a continuous stream of raw accelerometer and gyroscope data from which to extract the current path. It re-samples the input data to 20Hz before storing it into a memory buffer (2.3MB/h each with 32bit precision), which is a fixed circular buffer holding the past hour of measurements. The values are classified on-demand with LR and the classification results are stored into a separate circular buffer (3.6kB/h with 8 bit booleans). When calculating gravity and linear accelerometer values from raw accelerometer data,

TRec's memory buffer requirements are brought down to less than 5MB in total. We use weka[13] to model logistic regression on Android. The data is only smoothed when trajectory comparison is needed. Once triggered, comparison is done continuously at one second intervals for a set of ten attempts by default. If successful, the data is written to disk and a response is generated. Otherwise the authentication attempt is aborted.

IV. Evaluation

A. Movement Recognition

To evaluate the accuracy of movement recognition, we collected a preliminary dataset covering different motion conditions over four hours. In order to reduce potential overlearning, we regularized the training data by applying Gaussian additive and multiplicative noise (with $std = 0.05$ and $std = 0.02$). We then evaluated the resulting logistic regression (LR) model using five-fold stratified cross-validation to separate training and testing data.

Out of ten features, we found that the three most significant features were the standard deviation of the 5 second (and 1 second) sliding window of 3D differenced accelerometer values, and the peak-to-peak value for the 5 second sliding window of gyroscope measurements. These features corresponded to 58% of the weights of the LR model.

Evaluation results for LR showed a true positive rate of 98% for movement (M), and 92% for stationary (S). To improve the classification accuracy on out-of-sample data we additionally employ a Hidden Markov Model (HMM), with a prior assumption that state changes occur seldomly. We use the calculated misclassification probabilities as emission probabilities in HMM and use a default value of 99% probability to switch between hidden states (transition probabilities).

B. Experimental Data Acquisition

To evaluate the accuracy of TRec as a whole, we collected real-world usage data by repeatedly traversing a series of routes. The dataset consists of paths corresponding to 20 different, 6- to 12-minute-long routes in Espoo and Oxford. In order to generalize across different devices, we gathered data from five different device models[3] at 200Hz sampling rate. To ensure that gravity estimates are reliable, we attached the measurement device firmly to bicycles and pushed along. Each trajectory was repeated between 7 and 11 times. Routes contain real-world obstacles, such as traffic lights, gravel, asphalt or cobblestone roads and crowds. In total, the 7.7 GB dataset consists of the equivalent of 38.6 hours of recording, collected over a total distance of 123 km.

C. Path Similarity Thresholds

Separability. For each of 20 unique routes $R_k, k \in \{1, \ldots, 20\}$, our data contains between 7 and 11 instances i_{R_k}. In order to evaluate the within-class similarities of each route and determine if the Needleman-Wunsch measure supports consistent classification, we calculated the similarities between all pairs of instances, trimmed to 2-minute duration. For each route, this results in at least 21 unique pairwise within-class similarities, and 1595 unique pairwise between-class similarities.

Figure 7 shows an example of within-class and between-class similarities for one such route in our dataset, with all other routes showing similar behavior. In an ideal noiseless case, all instances i_R of a route R are identical, but in real-world data, sensor noise results in a a spread of similarities. Within-class similarities are spread in the range [18,32]. While there is an overlap of within-class and between-class similarities in the low range [18,21], most within-class and between-class cases are separated, implying that Needleman-Wunsch is indeed a good measure of similarity.

Determining the Initial threshold. As described in Section III, decision thresholds that TRec uses are adapted based on the number of instances of a reference path seen so far. When only a single instance has been authorized by the user, TRec uses the *initial threshold*; as more instances are seen, their similarities to the reference path are used to compute a *local* threshold, and subsequently the *mixed* threshold, which is a combination of initial and local thresholds that depends increasingly less on the initial threshold as the number of seen instances of a reference path increases.

We can compute the threshold with respect to a chosen trade-off between FAR and FRR by minimizing the combined error rate $\alpha \cdot FRR + (1 - \alpha) \cdot FAR$ for a specific value of the trade-off parameter α. Figure 8 shows optimal achievable FARs and FRRs for pooled[4] reference paths. For each value of α, the minimum combined FAR/FRR is found by varying the decision threshold in accordance to the resulting evaluation function. In scenarios where the usability of the system is more important, larger α values should be used.

The decision threshold is naturally tied to the length L of the reference path. In order for TRec to make a classification decision for a candidate path given only one instance of a reference path, we need to determine a function for the initial threshold that depends on the length of the reference path L. We did this by searching for the value d_L that minimizes the combined error rate for $\alpha \in \{0.1, 0.2, \ldots, 0.9\}$ for reference path lengths $L \in \{1, 2, \ldots, 6\}$. This gave us 6 dependent-independent variable pairs per α, i.e. 9 regressions. We found that the relationship between optimal decision threshold and the length of the path is affine. Initial decision thresholds are obtained by fixing the α value and using linear regression [19] with the temporal length L as the independent variable and the decision threshold d_L as the dependent variable.

Figure 8 shows that $\alpha^* = 0.5$ is closest to the equal error rate marker. Setting α^* to 0.5, we obtained the decision thresholds for pooled scores as:

$$D^*(L) = 9.69L - 1.40 \tag{3}$$

$D^*(L)$ is the function representation that estimates the previously introduced d_L.This rounds to $D^*(1) = 8$, $D^*(2) = 18$, $D^*(3) = 28$, $D^*(4) = 37$, $D^*(5) = 47$ and $D^*(6) = 57$. These particular decision thresholds serve as examples on how

[3]OnePlus One, Nexus 6, Moto G (gen 2), Nexus 5X and Samsung GS6.

[4]We aggregated all reference-path specific within-class distances to one group, and all reference-path specific between-class distances to another group.

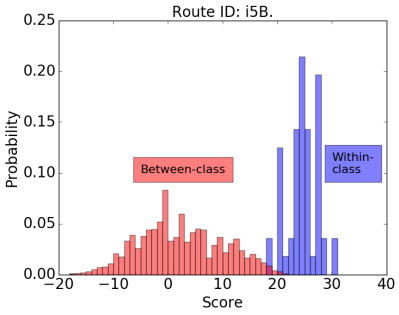

Fig. 7. An example of within-class (blue) and between-class similarities (red) for one route. Choosing a reference path-specific threshold 18 minimizes equal error rate (EER) on this route.

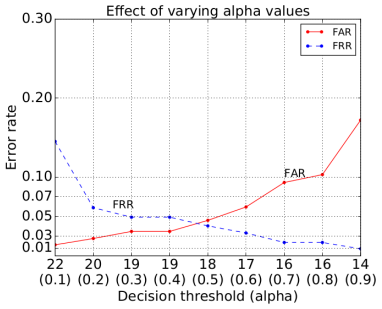

Fig. 8. Error rates vs. α (with corresponding decision threshold). Large α values result in high FRR, while small α results in high FAR.

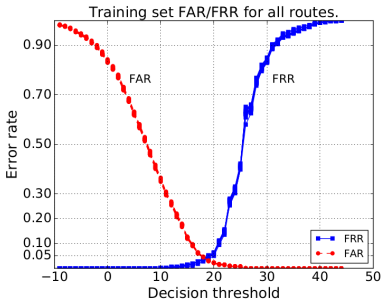

Fig. 9. Training set FAR (red) and FRR (blue) for 20 routes (20 graphs). Training set for each route consists of paths of *all other routes* pooled together. EER is achieved at 18 in all graphs.

TABLE I. MEAN AND STANDARD DEVIATION OF FAR/FRR USING ONLY THE **INITIAL DECISION THRESHOLD** INDIVIDUALLY FOR 20 PATHS.

L (min)	FAR	FRR
1.0 min	**0.064** $\pm$ 0.057	**0.137** $\pm$ 0.143
2.0 min	**0.043** $\pm$ 0.048	**0.039** $\pm$ 0.102
3.0 min	**0.073** $\pm$ 0.085	**0.047** $\pm$ 0.107
4.0 min	**0.072** $\pm$ 0.074	**0.061** $\pm$ 0.133
5.0 min	**0.049** $\pm$ 0.061	**0.057** $\pm$ 0.103
6.0 min	**0.047** $\pm$ 0.070	**0.040** $\pm$ 0.081

initial thresholds are calculated, and we use them in our `TRec` prototype. However, in order to give an unbiased estimate of system performance, we do not use exactly these values in the remainder of system evaluation, as this might lead to over-fitting and reporting overly optimistic results. Table V (Appendix) shows decision thresholds for pooled scores, using different values of the trade-off parameter α.

Instead, we calculate decision thresholds for each reference path separately, using the *leave-one-reference-path-out* method, which we use for all subsequent analysis (a training set is constructed using 19 reference paths; it is used to determine the initial threshold for the remaining reference path). Figure 9 shows the FARs and FRRs by removing one and retaining the 19 other reference paths (there are 20 lines in total). As can be seen in Figure 9, training set error rates below 5% are achievable for all 20 sets.

The next two subsections discuss actual test error rates for specific paths that use individual and mixed thresholds, evaluated w.r.t. two different parameters: the reference path length and the number of instances of a reference path.

D. Impact of Reference Path Length

The mean error rates and standard deviations for our 20 reference paths using the initial threshold with varying reference path lengths are shown in Table I. FRRs are calculated over all unique combinations of dividing instances i_R of routes R into reference path instances (training) and candidate paths (testing). FARs are calculated by treating all other instances $i_{\bar{R}}$ of routes $\bar{R}$ as candidate paths. We know the temporal length of the reference path and require the same length for each candidate path. Both FAR and FRR drop when increasing the length of the path from one to two minute, and reach lowest mean FAR and FRR on 2 and 6 min long reference paths. Initially, the mean FAR is between **4.3%** and **7.3%** and FRR between **3.9%** and **6.1%** for reference paths longer than two minutes. However, the variation in FRRs is high between different reference paths. We see a clear drop at 2 min, but FARs/FRRs for individual paths do not change significantly beyond $L = 2$ min (Wilcoxon signed rank test [17][5]). We thus use $L = 2$ min for the rest of the analysis in this paper.

E. Using Multiple Reference Path Instances

Multiple instances of a reference path can be used to increase security and usability of the system. This is done by selecting a good representative for the reference path among the seen instances (the instance at the cluster center - the medoid) and by adopting a mixed threshold using equations 1 and 2 (combination of d_i and d_l). The local decision threshold is calculated with the same scheme as the initial decision threshold, by setting $\alpha = 0.5$. If there are multiple solutions that provide the same minimum error rate, we use the decision threshold that is center most among these (maximum margin). The difference to the initial threshold derivation is that we generate paths to estimate the distribution of between-class similarities using a Markov chain [19]. This is done to prevent over-learning by using collected, real world paths. The Markov chain represents the movement primitive transition probabilities that $\mathcal{P}$ has experienced so far, i.e. probabilities $p(c \rightarrow d)$, where $c, d \in \{M, L, R\}$. It is defined with these 9 parameter values.

To evaluate the effects of using the medoid and Equation 2 to calculate the confidence factor λ we plotted the median FARs and FRRs (per reference path), by varying the number of instances and the decision threshold calculation scheme. It is easier to investigate trends with the median, since it is a robust estimator that tolerates outliers better than the mean. Figure 10 shows the median FARs and FRRs for the initial, local and mixed thresholds using 2 min reference paths. These represent doing no update, doing a full update or a conservative update to the decision threshold, respectively. At each update, the pairwise similarities are calculated between the instances, and the instance with the largest summed similarity to the other instances is selected as the medoid. The medoid is used for calculating the similarity to each new candidate path. The

[5]Null hypothesis: No change in median FAR (resp. FRR) by increasing paths' temporal length. FAR: p-value 0.47, FRR: p-value 0.64. Wilcox zero treatment, sample size = 79, test statistics 182.5 and 1023.

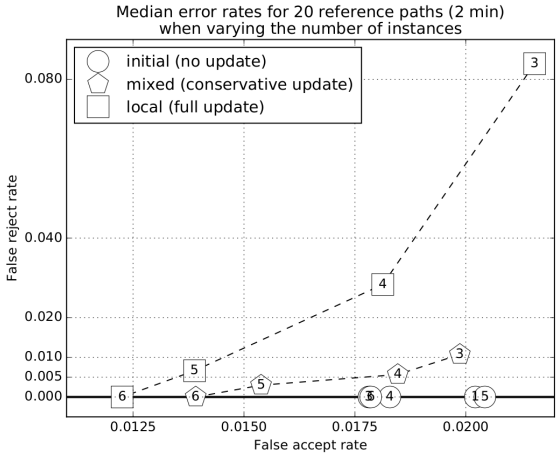

Fig. 10. Effect of increasing **number of instances** of 2 min reference paths, using different decision threshold update schemes. Coordinates describe median (FAR, FRR), numbers describe number of instances used.

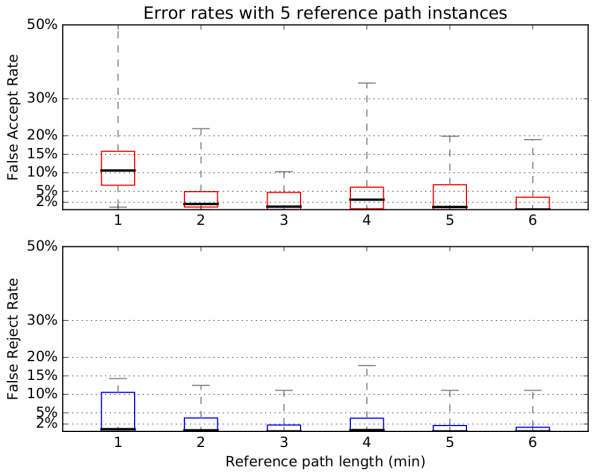

Fig. 11. **Mixed decision thresholds**. Box-and-whiskers plot over FARs and FRRs of different temporal lengths when TRec has seen **five instances** of a reference path. On average FARs are below 3% and FRRs are below 0.5% when $L \geq 2min$.

benefit in using the medoid to represent the reference path is that only one similarity score needs to be calculated, and that the score is calculated on the most representative instance.

In Figure 10, circles represent the continued usage of the initial threshold. The median FRR is zero, but the FAR does not improve with more instances, because the FAR depends on the decision threshold, not on the reference path representative. Using the initial threshold is equivalent to a constant confidence factor $\lambda = 0$, (no trust in d_l).

The squares represents the effect of using only the local threshold. While FARs are at a similar level as the initial threshold, median FRRs are significantly worse. With five or more instances, median FARs drop to a lower level than with initial thresholds. Using only the local threshold is equivalent to a constant confidence factor $\lambda = 1$ (full trust in d_l).

Mixed thresholds are derived from initial and local threshold values using equations 1 and 2, and are shown with pentagons. The mixed thresholds retain the good FRR of initial thresholds when few reference path instances are seen, and achieve improved performance similar to local thresholds when more instances have been observed. Using the confidence factor in conjunction with medoids to calculate decision thresholds is empirically shown to increase performance on two-minute paths, dropping the median FAR from $2.0\% \rightarrow 1.5\%$ when increasing the number of instances from 1 to 5, while the median FRR increases from $0.0\% \rightarrow 0.3\%$. Simultaneously, the mean FAR drops from $4.3\% \rightarrow 3.4\%$ and mean FRR drops from $3.8\% \rightarrow 2.8\%$. To further validate our results, we analyze the results of TRec when **five instances** are used with **mixed thresholds**. Table II shows the resulting mean FARs and FRRs for different reference path lengths. Mean FARs are in the range of **2.3%** to **5.0%** and mean FRRs are in the range of **1.8%** to **3.1%** for reference paths longer than 2 min. Both FAR and FRR are smaller for these paths, and the spreads in FRRs are significantly lower than in Table I. We find again that FARs/FRRs for individual paths do not change significantly by considering reference paths beyond 2 min (Wilcoxon signed rank test[6]).

TABLE II. MEAN AND STANDARD DEVIATION OF FAR/FRR USING **FIVE INSTANCES** AND **MIXED DECISION THRESHOLDS** FOR 20 PATHS.

L (min)	FAR	FRR
1.0 min	**0.141** $\pm$ 0.135	**0.046** $\pm$ 0.054
2.0 min	**0.034** $\pm$ 0.049	**0.028** $\pm$ 0.042
3.0 min	**0.023** $\pm$ 0.028	**0.018** $\pm$ 0.031
4.0 min	**0.050** $\pm$ 0.078	**0.031** $\pm$ 0.050
5.0 min	**0.042** $\pm$ 0.058	**0.017** $\pm$ 0.032
6.0 min	**0.030** $\pm$ 0.050	**0.015** $\pm$ 0.032

Figure 11 shows box-and-whiskers plots [19] for FARs and FRRs for our 20 reference paths, evaluated over different lengths, once five instances have been observed by TRec. The outer lines of the boxes denote the 75th and 25th percentiles, while the line in the boxes denotes the median. The whiskers denote the range, i.e. maximum and minimum values. Note that on average, FARs and FRRs are lower than the mean values reported in Table II. For reference paths longer than 2 min, median FARs are below 3% and FRRs below 0.5%.

Our implementation of TRec uses mixed thresholds so that its performance improves with time.

F. Energy Consumption

To evaluate requirement R3, we created a controlled experiment to measure the energy consumption of TRec. We obtained 3-hour consumption reports on three different devices[7], ten times each. A control group had no TRec installed, and did not have WiFi nor mobile connectivity.

Case 1. Once every hour, $\mathcal{U}$ walks 5 min and returns to $\mathcal{V}$. Walking initiates data collection at full sampling rate. Upon

[6]Null hypothesis: No change in median FAR (resp. FRR) by increasing paths' temporal length. FAR: p-value 0.43, FRR: p-value 0.15. Wilcox zero treatment, sample size = 79, test statistics 779.5 and 272.

[7]Nexus 5, Nexus 5X and Samsung GS6. The devices are charged to 100% before the start of the experiment; the battery level is measured after 3 hours.

approaching $\mathcal{V}$, path recognition is triggered and runs 60 times. This is intended to model usage by an office worker who works at a desk but takes a short break periodically to walk about. Table III reports the net contribution of TRec, compared to the control group.[8] The increase in energy consumption varies between devices. On average we recorded a drain of 0.56% (of total battery capacity) per hour.

TABLE III. NET INCREASES IN HOURLY BATTERY CONSUMPTION RATES (PERCENTAGE OF TOTAL CAPACITY) AND MILLIAMPERE (CASE 1).

Device	$\%/h$	mA
Nexus 5	0.41 ± 0.23	(9.32 ± 5.22)
Nexus 5X	0.63 ± 0.33	(17.13 ± 8.92)
Samsung S6	0.61 ± 0.24	(15.45 ± 6.44)

Case 2. We observed that the largest contributor to battery usage was the wakelock that ensured sensor measurements were processed quickly and not lost due to the sensor buffer filling up. Therefore, we recorded the energy usage with wakelock acquired for three hours. This represents the worst case scenario. The hourly net battery drains for our devices are shown in Table IV. The mean increase in battery usage varies across devices, with an overall average of 1.9% per hour. We believe that the drain increase in Nexus 5X and Samsung S6 is a result of these devices using different power states efficiently, causing the consumption in the control group to be lower.

TABLE IV. NET INCREASES IN HOURLY BATTERY CONSUMPTION RATES (PERCENTAGE OF TOTAL CAPACITY) AND MILLIAMPERE (CASE 2).

Device	$\%/h$	mA
Nexus 5	0.87 ± 0.28	(20.11 ± 6.40)
Nexus 5X	2.31 ± 0.24	(62.28 ± 6.60)
Samsung S6	2.55 ± 0.17	(65.14 ± 4.33)

Case 3. Energy consumption of settings and sensors can vary significantly between different smartphone models. On average users recharge their batteries within an hour of depletion [30]. A study [22] found the average battery life to be nine hours (11% consumption rate) under favorable circumstances. In a setting where $\mathcal{U}$ commutes for one hour and does office work for eight hours, we estimate (using data from Tables III and IV) the additional battery drain due to TRec to be between 4% and 7%, depending on the phone model. These correspond to a decrease in battery life of approximately 20-40 minutes for a device with a nine hours of battery life. For comparison, the loss in battery life when moving from an area with average WiFi to bad coverage was estimated to be 6.29% [22].

V. SECURITY ANALYSIS

TRec prevents relay attacks by having $\mathcal{P}$ verify its proximity to $\mathcal{V}$ through prover-side trajectory comparison: if $\mathcal{P}$'s current trajectory (candidate path) matches the authorized approach trajectory (reference path) to $\mathcal{V}$, $\mathcal{P}$ locally decides to participate in the authentication protocol with $\mathcal{V}$.

To evaluate the security guarantees of TRec, we analyze our dataset (Section IV) corresponding to a total of 123 km equivalent to 38.6 hours of movement. We used false accept rate (FAR) of trajectory comparison as the measure of security. We showed that the adversary $\mathcal{A}$ has less than 5% chance of success when TRec has seen only one instance of the reference path (Table I). The success rate steadily drops over time, as more instances are seen by TRec, dropping below 3.5% on average (Table II) and ∼1.5% median (Figure 10) after 5 instances. Note that to achieve such rates $\mathcal{A}$ is required to follow $\mathcal{U}$ in close proximity while $\mathcal{U}$ is in motion carrying $\mathcal{P}$. The typical relay attack scenario, e.g., a thief stealing a keyless-entry car, is where $\mathcal{P}$ is stationary [33]. TRec *completely prevents relay attacks in this scenario* as no candidate path is generated by $\mathcal{P}$. As such, TRec substantially improves security over existing schemes as it makes relay attacks significantly harder.

The complexity of the trajectory leading to $\mathcal{V}$ is inversely proportional to the likelihood that there exists another route (not terminating in $\mathcal{V}$) that $\mathcal{A}$ could use to mount a relay attack against TRec *if* $\mathcal{U}$ happens to traverse it. The measure for complexity however depends on the geographic neighborhood: for instance, a trajectory with many turns can be considered to be more complex than a straight-line trajectory in most cities but in an old town that has few straight roads, a straight-line trajectory may be unique. We can incorporate such uniqueness assessment of reference path into the TRec app by analyzing a city's map data to find paths that are too similar to a given reference path and hence vulnerable to relay attacks. This provides the user with an estimate of the likelihood of TRec unintentionally "unlocking" $\mathcal{V}$ while the user is traversing some other route not leading to $\mathcal{V}$. TRec already offers the option for the user to choose the length of the reference path (Figure 6).

VI. DISCUSSION AND FUTURE WORK

Why transparent authentication. When trajectory comparison fails, TRec will prompt the user for explicit confirmation as to whether the current candidate path should be included in the set of legitimate reference paths for a given verifier. A natural question at this point is whether the complexity of TRec could have been avoided by simply dispensing with the transparency requirement and asking users to always manually confirm proximity. However, the trend towards transparent authentication is unmistakable. Google's Trust API (Project Abacus) [15] underscores this trend. The popularity of keyless entry as a premium feature across different car manufacturers suggest that consumers are willing to even pay for the convenience. TRec's contribution is to *retain* the convenience promised by transparent authentication while significantly enhancing security against relay attacks.

Verification of Return Paths. By relying on reference path instances gathered over time, TRec is currently limited to scenarios with "fixed" verifiers. But since TRec represents trajectories as sequences of primitives it is possible to compare a reference path that that *starts* from the location of the verifier to a candidate path that ends at the verifier by reversing the reference path. This allows transparent authentication even for *mobile verifiers*: e.g., when a user parks his car at a new place and walks to a mall and subsequently returns to the car via the same route. Classification performance relies primarily on the reproducibility of path sequences; the direction of movement should not impact the generated sequences. Our informal evaluation of this approach showed promising results. But we leave its comprehensive analysis for future work.

[8] Android Debug Bridge provides lower and upper bounds to the battery usage with 1% granularity. We only show the upper bound estimates here.

Prover Rotation. If $\mathcal{P}$ is a portable device like a smartphone, fast changes in $\mathcal{P}$'s perceived gravity direction (e.g. taking the phone out of the pocket) affect turn detection reliability. In our evaluation, we avoided this by affixing $\mathcal{P}$ firmly. We verified that fixing $\mathcal{P}$'s position to the user's torso[9] produced similar reliable results as fixing $\mathcal{P}$ to a bicycle, as the gravity estimate is reliable in both cases. Our results apply in cases where $\mathcal{P}$ is a wearable device integrated into torso clothing or an accessory integrated into a vehicle like a bicycle or car.

Additional Sequence Primitives. While this paper focuses only on using low-cost internal sensors, `TRec` can be extended to use additional primitives besides only movement/left/right. If the requirement to use only internal sensors is relaxed, the system could start detecting presence of specific wireless networks or indoor short-range beacons; if the user is indeed taking the same path as was the case when the corresponding reference path was recorded, then the same events should be detected at specific locations. Such detected presences can easily be represented by unique symbols and seamlessly integrated into the current implementation of `TRec`. Additionally, while we currently focus only on a single movement modality, `TRec` can be augmented with a transport mode detection scheme [14] that would add further entropy to generated paths. As such, a path that includes walking, then taking a bus, and then walking would be additionally more specific than if all movement intervals are represented with the same symbol.

VII. Related Work

Trajectory Recognition. Trajectory recognition estimates the location of a user and can be used to enable additional convenience functions but also to subvert a user's privacy. The following techniques do not rely on GPS location information but use a different way to obtain location ground truth, e.g. a street map or the timetable of public transportation. After a correlation between mobile measurements and this information the system estimates the user's location.

Gao et al. [8] propose an approach using vehicular speed and the start location to estimate final destination and path of a car. Therefore, speed measurements are recorded and segmented assuming stops at traffic lights and deceleration throughout turns. By matching these segments to map data they achieved an accuracy of 500 meters for 24% of traces in the New Jersey and 26% of traces in the Seattle area. The low accuracy of this approach makes it unsuitable for a security critical application.

Further, Watanabe et al. [31] identify a user's train trips based on inertial measurements. First, the user's activity is classified into inside a *vehicle*, *walking* and remaining *stationary*. Afterwards, the transition times between the different modes are used to correlate them with timetables. Each train trip is weighted according to its popularity to reduce the number of candidates. Their results show that a location detection is feasible although limited to paths along the train network which makes it infeasible for co-location verification.

The work by Nawaz and Mascolo [20] explores the significant transport routes of a user based on gyroscope data. According to their hypothesis, a route exposes a certain signature based on angular momentum. Due to traffic conditions, however, the different instances of the same route vary in length. To account for this, they apply dynamic time warping (DTW). The DTW scores are then used to form clusters of observed paths. Our system in contrast ignores stationary phases which makes time warping unnecessary. The co-presence is verified by comparing string representations of the paths.

Co-presence verification. The co-location of devices is an important countermeasure against impersonation and relay attacks. In some cases it is also used as a second factor for authentication. Although GPS could be used to assert co-location in theory, these signals are not authenticated and thus not trustworthy [27]. A range of alternatives based on context comparison has recently emerged.

Halevi et al. [12] proposed co-presence detection based on comparing audio and light. A merchant terminal and mobile phone probe their environments to compare them to assert co-location. They evaluate both modalities separately, and achieve a FAR of 6.5% and a 5% FRR for light while reporting a FAR and FRR of 0% for audio. The experiment is carried out using two identical devices at immediate proximity, which makes the result overly optimistic.

A similar approach to mitigate relay attacks was explored by Shrestha et al. [23] and Truong et al. [28]. They use natural environment properties as well as digital signals. Truong et al. [28] identify WiFi as the dominating feature with a FAR of 2% and a FRR of 1%. In their approach, Shrestha et al. conclude that a modality fusion reduces the FRR of up to 24% and FAR of up to 33% of an individual features to 3% and 6% respectively. However, in follow-up work they were able to increase the FAR from 3% to 66% by manipulating a single modality [24]. Hence, increasing the number of modalities does not necessarily strengthen security as it depends on the weights machine learning models assign to them. A thorough analysis of these algorithms is required to give sophisticated security guarantees.

Karapanos et al. [18] use the audio fingerprint of a location as a second factor for authentication. They implemented a system that uses a browser and a co-located smart phone to record environmental noise for comparison. Their threat model assumes a remote attacker who obtained the user's credentials. Although, this approach seems promising against relay attacks their use case and evaluation do not accommodate for it.

VIII. Conclusion

In this paper we proposed, built, and evaluated `TRec` as a novel approach to prevent relay attacks in signal-strength-based transparent authentication schemes. `TRec` is based on the observation that device location can be verified from its movement history. As a result, the prover device only participates in the authentication protocol if it can verify its proximity to the verifier. This is done by comparing the current trajectory to authorized trajectories ending at the specific verifier. `TRec` measures trajectories using low-cost sensors that are commonly available on personal devices, and does not rely on external signals that can be under adversary's control.

We demonstrate the security of `TRec` against an adversary who follows the user and relays wireless communication.

[9] Either held in hand or fixed to clothes.

By acquiring and analyzing empirical data, collected using a TRec prototype, we show that TRec's prover-side proximity verification efficiently prevents relay attacks, with low false rejects. Our analysis shows that the scheme achieves good security / usability trade-offs and is energy efficient.

TRec can be realized entirely within the prover device. Thus it can be easily integrated into *existing* transparent authentication schemes, which significantly lowers the barrier for deployment of TRec.

ACKNOWLEDGMENTS

This work was supported by the Academy of Finland "Contextual Security" project (274951) and the "CDT Cyber Security" fund of the Engineering and Physical Sciences Research Council, United Kingdom.

REFERENCES

[1] Android. Keeping the device awake, 2016. https://developer.android.com/training/scheduling/wakelock.html.

[2] Joseph Bonneau, Cormac Herley, Paul C. van Oorschot, and Frank Stajano. The quest to replace passwords: A framework for comparative evaluation of web authentication schemes. In *IEEE Symposium on Security and Privacy, SP 2012, 21-23 May 2012, San Francisco, California, USA*, pages 553–567. IEEE Computer Society, 2012.

[3] Joseph Bonneau, Cormac Herley, Paul C. van Oorschot, and Frank Stajano. Passwords and the evolution of imperfect authentication. *Commun. ACM*, 58(7):78–87, 2015.

[4] Stefan Brands and David Chaum. Distance-Bounding Protocols. In *Advances in Cryptology EUROCRYPT '93*, pages 344–359. Springer Berlin Heidelberg, Berlin, Heidelberg, 1993.

[5] Mark D. Corner and Brian D. Noble. Zero-interaction authentication. In *Proceedings of the 8th annual international conference on Mobile computing and networking - MobiCom '02*, page 1, New York, New York, USA, sep 2002. ACM Press.

[6] Danny Dolev and Andrew C Yao. On the security of public key protocols. *Information Theory, IEEE Transactions on*, 29(2):198–208, 1983.

[7] Richard Durbin, Sean R Eddy, Anders Krogh, and Graeme Mitchison. *Biological sequence analysis: probabilistic models of proteins and nucleic acids*. Cambridge university press, 1998.

[8] Bernhard Firner, Shridatt Sugrim, Yulong Yang, and Janne Lindqvist. Elastic pathing: Your speed is enough to track you. *arXiv preprint arXiv:1401.0052*, pages 975–986, 2013.

[9] Aurelien Francillon, Boris Danev, and Srdjan Capkun. Relay Attacks on Passive Keyless Entry and Start Systems in Modern Cars. *Network and Distributed System Security Symposium*, pages 431–439, 2011.

[10] Lishoy Francis, Gerhard Hancke, Keith Mayes, and Konstantinos Markantonakis. Practical relay attack on contactless transactions by using NFC mobile phones. In *Cryptology and Information Security Series*, volume 8, pages 21–32, 2012.

[11] Tzipora Halevi, Di Ma, Nitesh Saxena, and Tuo Xiang. Secure proximity detection for nfc devices based on ambient sensor data. In *European Symposium on Research in Computer Security*, pages 379–396. Springer, 2012.

[12] Tzipora Halevi, Di Ma, Nitesh Saxena, and Tuo Xiang. Secure Proximity Detection for NFC Devices Based on Ambient Sensor Data. pages 379–396. Springer Berlin Heidelberg, 2012.

[13] Mark Hall, Eibe Frank, Geoffrey Holmes, Bernhard Pfahringer, Peter Reutemann, and Ian H Witten. The weka data mining software: an update. *ACM SIGKDD explorations newsletter*, 11(1):10–18, 2009.

[14] Samuli Hemminki, Petteri Nurmi, and Sasu Tarkoma. Accelerometer-based transportation mode detection on smartphones. In *Proceedings of the 11th ACM Conference on Embedded Networked Sensor Systems*, page 13. ACM, 2013.

[15] Alex Hern. Google aims to kill passwords by the end of this year — Technology — The Guardian, 2016.

[16] Antonio R Jimenez, Fernando Seco, Carlos Prieto, and Jorge Guevara. A comparison of pedestrian dead-reckoning algorithms using a low-cost mems imu. In *Intelligent Signal Processing, 2009. WISP 2009. IEEE International Symposium on*, pages 37–42. IEEE, 2009.

[17] Eric Jones, Travis Oliphant, P Peterson, et al. SciPy: Open source scientific tools for python, 2001.

[18] Nikolaos Karapanos, Claudio Marforio, Claudio Soriente, and Srdjan Capkun. Sound-Proof: Usable Two-Factor Authentication Based on Ambient Sound. In *24th USENIX Security Symposium (USENIX Security 15)*, pages 483–498, 2015.

[19] Kevin P Murphy. *Machine learning: a probabilistic perspective*. MIT press, 2012.

[20] Sarfraz Nawaz and Cecilia Mascolo. Mining users' significant driving routes with low-power sensors. In *Proceedings of the 12th ACM Conference on Embedded Network Sensor Systems*, pages 236–250. ACM, 2014.

[21] Vishal M. Patel, Rama Chellappa, Deepak Chandra, and Brandon Barbello. Continuous User Authentication on Mobile Devices: Recent progress and remaining challenges. *IEEE Signal Processing Magazine*, 33(4):49–61, jul 2016.

[22] Ella Peltonen, Eemil Lagerspetz, Petteri Nurmi, and Sasu Tarkoma. Energy modeling of system settings: A crowdsourced approach. In *Pervasive Computing and Communications (PerCom), 2015 IEEE International Conference on*, pages 37–45. IEEE, 2015.

[23] Babins Shrestha, Nitesh Saxena, Hien Thi Thu Truong, and N. Asokan. Drone to the Rescue: Relay-Resilient Authentication using Ambient Multi-sensing. pages 349–364. Springer Berlin Heidelberg, 2014.

[24] Babins Shrestha, Nitesh Saxena, Hien Thi Thu Truong, and N. Asokan. Contextual Proximity Detection in the Face of Context-Manipulating Adversaries. Nov 2015. arXiv report /1511.00905 http://arxiv.org/abs/1511.00905.

[25] SourceForge. BlueProximity. https://sourceforge.net/projects/blueproximity/.

[26] Anastasios Tefas, Constantine Kotropoulos, and Loannis Pitas. Using Support Vector Machines to enhance the performance of elastic graph matching for frontal face authentication. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 23(7):735–746, 2001.

[27] Nils Ole Tippenhauer, Christina Pöpper, Kasper Bonne Rasmussen, and Srdjan Capkun. On the requirements for successful GPS spoofing attacks. *ACM conference on Computer and communications security, CCS*, page 75, 2011.

[28] Hien Thi Thu Truong, Xiang Gao, Babins Shrestha, Nitesh Saxena, N. Asokan, and Petteri Nurmi. Using contextual co-presence to strengthen Zero-Interaction Authentication: Design, integration and usability. *Pervasive and Mobile Computing*, 16:187–204, 2015.

[29] Hien Thi Thu Truong, Babins Shrestha, Nitesh Saxena, N. Asokan, and Petteri Nurmi. Comparing and fusing different sensor modalities for relay attack resistance in Zero-Interaction Authentication. In *2014 IEEE International Conference on Pervasive Computing and Communications (PerCom)*, pages 163–171. IEEE, mar 2014.

[30] Daniel T Wagner, Andrew Rice, and Alastair R Beresford. Device analyzer: Understanding smartphone usage. In *International Conference on Mobile and Ubiquitous Systems: Computing, Networking, and Services*, pages 195–208. Springer, 2013.

[31] Takuya Watanabe, Mitsuaki Akiyama, and Tatsuya Mori. RouteDetector: Sensor-based Positioning System That Exploits Spatio-Temporal Regularity of Human Mobility. 2015.

[32] Grady Xiao, Mariofanna Milanova, and Mengjun Xie. Secure behavioral biometric authentication with leap motion. In *2016 4th International Symposium on Digital Forensic and Security (ISDFS)*, pages 112–118. IEEE, apr 2016.

[33] Yle News. Finnish police: Keep your car keys in the fridge. Yle News report, Sep 2016. http://yle.fi/uutiset/osasto/news/finnish_police_keep_your_car_keys_in_the_fridge/9166149.

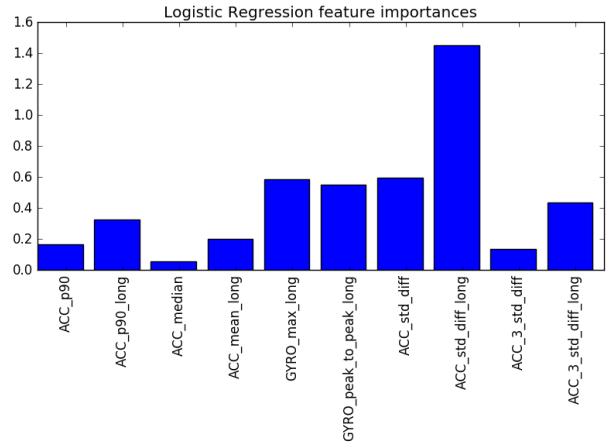

Fig. 12. Feature importances (absolute weights) of the logistic regression model for movement classification. The sliding window standard deviation for five and one seconds were the most discriminative features.

APPENDIX

TABLE V. INITIAL DECISION THRESHOLDS FOR POOLED SCORES. THE R VALUE FOR THE FIT IS SHOWN IN PARENTHESES. L IS THE REFERENCE PATH LENGTH IN MINUTES.

α	$D(L)$
0.1	$10.171L + 1.400$ (r=0.998)
0.2	$10.314L - 0.600$ (r=0.999)
0.3	$10.143L - 1.000$ (r=0.999)
0.4	$9.543L - 0.067$ (r=0.998)
0.5	$9.686L - 1.400$ (r=0.998)
0.6	$9.914L - 2.533$ (r=0.998)
0.7	$9.771L - 3.533$ (r=0.996)
0.8	$8.600L - 1.600$ (r=0.993)
0.9	$7.914L - 2.200$ (r=0.989)

Table V shows initial decision thresholds obtained by minimizing the error $\alpha \cdot FRR + (1 - \alpha) \cdot FAR$. The trade-off variable α denotes how important initial FARs or FRRs are perceived to be, and may be application-specific. Higher α are more permissive (more weight on FRR), while lower α are more strict (less weight on FRR). The similarity between a candidate path and a reference path needs to be higher than the decision threshold in order to be accepted as valid instances of reference paths. Depending on how important FRRs or FARs are perceived to be, the initial decision thresholds may vary between 14[10] and 22[11] for a two-minute path.

Table VI shows the features extracted from segments of sensor readings. We used a feature selection algorithm to prune the feature set to ten features. These features are calculated every second, using either the past one or five seconds. Features 1 to 8 use the magnitude values obtained with the scalarization $m = \sqrt{x^2 + y^2 + z^2}$ of the three-dimensional linear accelerometer (or gyroscope) vectors. Features 9 and 10 perform operations on each axis separately and sum the results to one scalar value. Figure 12 shows the feature weights in the logistic regression algorithm.

Algorithm 1 shows a high-level scheme for calculating turns from angular data. It continuously estimates the sliding window standard deviation of the data and records turns when the standard deviation is above a threshold.

[10] $\alpha = 0.9$: $round(7.914 \cdot 2 - 2.200)$, with $L = 2$
[11] $\alpha = 0.1$: $round(10.171 \cdot 2 + 1.400)$, with $L = 2$

Algorithm 1 Turn detection with angular data.

while new event heading angle α_i detected **do**
  update *sliding window standard deviation* for last 2 seconds
  update estimate when turn *begins* and *ends*
  update *peak value* during turn
  if turn *ends* or *peak surpassed* **then**
    calculate how big the turn was in angles (d)
    calculate number of characters $n = |round(d/15°)|$
    if $d > 0$ **then**
      add n symbols R
    else[$d < 0$]
      add n symbols L

TABLE VI. FEATURES USED IN MOVEMENT RECOGNITION.

i	Feature name	Weight	Feature description
1	ACC_p90	-0.85	90th percentile of accelerometer magnitude value (1s)
2	ACC_p90_long	1.68	90th percentile of accelerometer magnitude value (5s)
3	ACC_median	0.29	Median of accelerometer magnitude (1s)
4	ACC_mean_long	-0.91	Mean of accelerometer magnitude (5s)
5	GYRO_max_long	2.96	Maximum of gyroscope magnitude (5s)
6	GYRO_peak_to_peak_long	2.81	(Maximum-minimum) for gyroscope magnitude (5s)
7	ACC_std_diff	-3.11	Standard deviation for differenced accelerometer magnitudes (1s)
8	ACC_std_diff_long	-7.19	Standard deviation for differenced accelerometer magnitudes (5s)
9	ACC_3_std_diff	-0.76	Sum of standard deviations for accelerometer values, separately differenced in each axis (1s)
10	ACC_3_std_diff_long	-2.29	Sum of standard deviations for accelerometer values, separately differenced in each axis (5s)
	Model intercept	3.09	Base assumption if other features are zero